

WHAT IS CLAIMED IS:

- 1 1. A method comprising:
 - 2 receiving a request to load a device policy module into a memory, wherein the
 - 3 device policy module is for use by a device driver, and wherein the
 - 4 device policy module includes at least one of a function, a procedure,
 - 5 and an object-oriented method operable to perform at least one of
 - 6 input/output (I/O) operation scheduling, path selection, and I/O
 - 7 operation error analysis;
- 8 loading the device policy module into the memory; and
- 9 informing the device driver of availability of the device policy module.

- 1 2. The method of claim 1 wherein the request to load a device policy module
2 into a memory is received from at least one of a user application and a device
3 discovery application.

- 1 3. The method of claim 1 wherein a portion of the memory comprises a kernel
2 memory space, and wherein the loading the device policy module into the memory
3 further comprises:
 - 4 loading the device policy module into the kernel memory space.

- 1 4. The method of claim 1 wherein the informing the device driver of
2 availability of the device policy module further comprises:
 - 3 registering the device policy module with the device driver by calling at least
 - 4 one of a function, a procedure, and an object-oriented method
 - 5 associated with the device driver.

- 1 5. The method of claim 1 further comprising:
 - 2 determining whether the device policy module is currently present in the
 - 3 memory.

- 1 6. The method of claim 1 further comprising:
 - 2 informing the device driver of unavailability of the device policy module.

- 1 7. The method of claim 6 wherein the informing the device driver of
2 unavailability of the device policy module further comprises:
3 unregistering the device policy module with the device driver by calling at
4 least one of a function, a procedure, and an object-oriented method
5 associated with the device driver.

- 1 8. The method of claim 1 wherein the device policy module is for use with a
2 corresponding storage device, the method further comprising:
3 transmitting at least one storage device attribute to the device driver.

- 1 9. The method of claim 1 wherein the at least one of a function, a procedure,
2 and an object-oriented method of the device policy module is specific to a particular
3 storage device.

- 1 10. The method of claim 1 wherein the at least one of a function, a procedure,
2 and an object-oriented method operable to perform at least one of I/O operation
3 scheduling, path selection, and I/O operation error analysis performs at least one of:
4 selecting one of a plurality of communication pathways to at least one storage
5 device;
6 selecting one or more sub-devices of the at least one storage device which will
7 be affected due to a communication pathway failure;
8 selecting an alternate communication pathway in case of a failure of one of the
9 plurality of communication pathways;
10 changing a current communications pathway from a first one of the plurality
11 of communication pathways to a second one of the plurality of
12 communication pathways;
13 responding to SCSI reservation/release requests; and
14 selectively transmitting I/O operations along at least two of the plurality of
15 communication pathways to the at least one storage device.

- 1 11. The method of claim 1 further comprising:
2 monitoring operation of the device policy module.

- 1 12. The method of claim 1 further comprising:
2 discovering the presence of at least one storage device belonging to a
3 distributed computing system.

- 1 13. The method of claim 12 further comprising:
2 determining whether the at least one storage device has a corresponding
3 device policy module.

- 1 14. A system comprising:
2 a storage device discovery module configured to determine information about
3 at least one storage device belonging to a distributed computing
4 system; and
5 a multipath driver in communication with the storage device discovery module
6 and configured to direct input/output (I/O) operations along at least one
7 of a plurality of communication pathways to the at least one storage
8 device, the multipath driver including:
9 an interface configured to communicate with a device policy module
10 including at least one of a function, a procedure, and an object-
11 oriented method operable to perform at least one of I/O
12 operation scheduling, path selection, and I/O operation error
13 analysis.

- 1 15. The system of claim 14 further comprising:
2 a device policy module including at least one of a function, a procedure, and
3 an object-oriented method operable to perform at least one of I/O
4 operation scheduling, path selection, and I/O operation error analysis.

- 1 16. The system of claim 15 wherein the at least one of a function, a
2 procedure, and an object-oriented method of the device policy module is specific to a
3 particular storage device.

- 1 17. The system of claim 14 wherein the at least one of a function, a
2 procedure, and an object-oriented method operable to perform at least one of I/O
3 operation scheduling, path selection, and I/O operation error analysis performs at least
4 one of:

5 select one of the plurality of communication pathways to the at least one
6 storage device;
7 select one or more sub-devices of the at least one storage device which will be
8 affected due to a communication pathway failure;
9 select an alternate communication pathway in case of a failure of one of the
10 plurality of communication pathways;
11 effect a communications pathway changeover;
12 respond to respond to SCSI reservation/release requests; and
13 selectively transmit I/O operations along at least two of the plurality of
14 communication pathways to the at least one storage device.

1 18. The system of claim 17 wherein the at least one storage device is a disk
2 array and wherein the one or more sub-devices are disk drives.

1 19. The system of claim 14 further comprising:
2 a memory; and
3 a processor coupled to the memory, wherein at least one of the storage device
4 discovery module and multipath driver are encoded as instructions
5 stored in the memory and executable on the processor.

1 20. The system of claim 19 wherein a first portion of the memory is used as a
2 kernel memory space and wherein a second portion of the memory is used as a user
3 memory space, and wherein the multipath driver is stored in the kernel memory space.

1 21. The system of claim 14 wherein the multipath driver further comprises:
2 a fixed set of I/O policies including at least one of a function, a procedure, and
3 an object-oriented method operable to perform at least one of I/O
4 operation scheduling, path selection, and I/O operation error analysis.

1 22. The system of claim 14 wherein the interface configured to communicate
2 with a device policy module includes at least one of a function, a procedure, and an
3 object-oriented method operable to perform at least one of registering a device policy
4 module with the multipath driver and unregistering a device policy module with the
5 multipath driver.

1 23. The system of claim 14 wherein the multipath driver is further configured
2 to monitor at least one loaded device policy module.

1 24. The system of claim 14 wherein the multipath driver is further configured
2 to receive at least one of a request to load a device policy module and a request to
3 unload a device policy module.

1 25. The system of claim 14 wherein the information about at least one storage
2 device includes at least one device attribute and wherein the device discovery module
3 is further configured to transmit the information about at least one storage device to
4 the multipath driver.

1 26. The system of claim 25 wherein the at least one device attribute includes
2 at least one of: a number of paths to the device, primary path information, secondary
3 path information, connected path information, disconnected path information, vendor
4 information, an enclosure serial number, and an LUN serial number, an array type.

1 27. The system of claim 14 wherein the storage device discovery module is
2 further configured to transmit the information about at least one storage device to the
3 multipath driver.

1 28. The system of claim 14 wherein the storage device discovery module is
2 further configured to receive at least one of a request to load a device policy module
3 and a request to unload a device policy module.

1 29. A computer readable medium comprising program instructions executable
2 on a processor, the computer readable medium being at least one of an electronic
3 storage medium, a magnetic storage medium, an optical storage medium, and a
4 communications medium conveying signals encoding the instructions, wherein the
5 program instructions are operable to implement each of:

6 receiving a request to load a device policy module into a memory, wherein the
7 device policy module is for use by a device driver, and wherein the
8 device policy module includes at least one of a function, a procedure,
9 and an object-oriented method operable to perform at least one of

10 input/output (I/O) operation scheduling, path selection, and I/O
11 operation error analysis;
12 loading the device policy module into the memory; and
13 registering the device policy module with the device driver.

1 30. The computer readable medium of claim 29 wherein the request to load a
2 device policy module into a memory is received from at least one of a user application
3 and a device discovery application.

1 31. The computer readable medium of claim 29 wherein a portion of the
2 memory comprises a kernel memory space, and wherein the program instructions
3 operable to implement the loading the device policy module into the memory further
4 comprise program instructions operable to implement:

5 loading the device policy module into the kernel memory space.

1 32. The computer readable medium of claim 29 wherein the program
2 instructions operable to implement the registering the device policy module with the
3 device driver further comprise program instructions operable to implement:
4 calling at least one of a function, a procedure, and an object-oriented method
5 associated with the device driver.

1 33. The computer readable medium of claim 29 further comprising program
2 instructions operable to implement:
3 determining whether the device policy module is currently present in the
4 memory.

1 34. The computer readable medium of claim 29 wherein the at least one of a
2 function, a procedure, and an object-oriented method of the device policy module is
3 specific to a particular storage device.

1 35. The computer readable medium of claim 29 wherein the at least one of a
2 function, a procedure, and an object-oriented method operable to perform at least one
3 of I/O operation scheduling, path selection, and I/O operation error analysis comprises
4 program instructions operable to perform at least one of:
5 selecting one of a plurality of communication pathways to at least one storage
6 device;

7 selecting one or more sub-devices of the at least one storage device which will
8 be affected due to a communication pathway failure;
9 selecting an alternate communication pathway in case of a failure of one of the
10 plurality of communication pathways;
11 changing a current communications pathway from a first one of the plurality
12 of communication pathways to a second one of the plurality of
13 communication pathways;
14 responding to SCSI reservation/release requests; and
15 selectively transmitting I/O operations along at least two of the plurality of
16 communication pathways to the at least one storage device.

1 36. The computer readable medium of claim 29 further comprising program
2 instructions operable to implement:
3 monitoring operation of the device policy module.

1 37. An apparatus comprising:
2 a means for directing input/output (I/O) operations along at least one of a
3 plurality of communication pathways to at least one storage device;
4 a means for providing storage device specific I/O operation scheduling and
5 communication pathway selection in conjunction with the means for
6 directing I/O operations; and
7 a means for selectively making the means for providing storage device
8 specific I/O operation scheduling and communication pathway
9 selection available to the means for directing I/O operations.

1 38. The apparatus of claim 37 further comprising:
2 a means for discovering the presence of at least one storage device belonging;
3 and
4 a means for communicating information about the at least one storage device
5 to the means for directing I/O operations.

1 39. The apparatus of claim 37 further comprising:
2 a means for registering the means for providing storage device specific I/O
3 operation scheduling and communication pathway selection with the
4 means for directing I/O operations.

1 40. The apparatus of claim 37 wherein the means for directing I/O operations
2 further comprises a means for providing a fixed set of I/O policies including at least
3 one of a function, a procedure, and an object-oriented method operable to perform at
4 least one of I/O operation scheduling, path selection, and I/O operation error analysis.